

Regional Appearance in Deformable Model Segmentation

Joshua V. Stough, Robert E. Broadhurst,
Stephen M. Pizer and Edward L. Chaney

Medical Image Display & Analysis Group (MIDAG)
University of North Carolina, Chapel Hill NC 27599, USA
`stough@cs.unc.edu`

Abstract. Automated medical image segmentation is a challenging task that benefits from the use of effective image appearance models. In this paper, we compare appearance models at three regional scales for statistically characterizing image intensity near object boundaries in the context of segmentation via deformable models. The three models capture appearance in the form of regional intensity quantile functions. These distribution-based regional image descriptors are amenable to Euclidean methods such as principal component analysis, which we use to build the statistical appearance models.

The first model uses two regions, the interior and exterior of the organ of interest. The second model accounts for exterior inhomogeneity by clustering on object-relative local intensity quantile functions to determine tissue-consistent regions relative to the organ boundary. The third model analyzes these image descriptors per geometrically defined local region.

To evaluate the three models, we present segmentation results on bladders and prostates in CT in the context of day-to-day adaptive radiotherapy for the treatment of prostate cancer. Results show improved segmentations with more local regions, probably because smaller regions better represent local inhomogeneity in the intensity distribution near the organ boundary.

1 Introduction

Automated segmentation of objects in medical images is a challenging task that benefits from the use of effective image appearance models. Some models that have shown success are those that analyze intensity patterns with respect to larger-than-voxel-scale regions of the image in the vicinity of and including the organ of interest. Such methods are compatible with the idea that the appearance of an organ in an image is in part a function of the position and intensities of neighboring organs and volumes and their tissue mixtures. In this paper we compare appearance models at three regional scales that statistically characterize image intensity near object boundaries for use in deformable model segmentation.




Fig. 1. Example sagittal CT slice of the male pelvis (left), with the bladder and prostate highlighted (right). Note the lack of contrast around the organs of interest.

In Bayesian deformable model segmentation, a geometric model for an object of interest is deformed via its shape parameters to fit the image data. The optimization proceeds according to an objective function that includes a term measuring geometry-to-image match, or “image match” for short, which is the result of the image appearance model. Some previous image matches depend on a voxel-scale dense correspondence across training cases. Such image matches include those based on intensity profiles that are either associated with individual image points [1] or averaged over similar image points [2], and those based on tuples that are formed from intensity-derived features at an ordered collection of these points [3]. These schemes are effective in situations where objects have a consistent voxel-scale structural relationship with one another. However, in segmenting organs in the male pelvis for example, one cannot expect the same tissue-mixture at a voxel-scale correspondence (see Fig. 1).

Region-based methods address this concern by modeling intensity distributions in object-relative regions. These approaches typically sample image intensities within the object interior or separately the interior and exterior. Some of the resulting intensity models use foreground/background intensity ranges [4], or use summary statistics such as mean and variance [5] [6]. Such simplifications of regional intensity distributions limit the information captured by the appearance model. Recently developed methods use an appropriate parameterization of the full intensity distribution for a region and either compute an image match with respect to a single reference distribution [7] [8], or directly model the variability in the intensity distribution as seen in training [9].

In [9] the authors sample from the interior and exterior regions and convert the resulting distributions to regional intensity quantile functions (RIQFs), which are amenable to linear statistical methods such as principal component analysis (PCA). Their appearance model is then based on probabilities of the regional distributions. As with previous approaches, the method has a significant drawback: the use of a single global exterior. Any such model oversimplifies the appearance by failing to account for the inhomogeneity in the local intensity distributions exterior to the object. While it may be reasonable to model intensities in the interior of an organ as samples from a single distribution, the exterior may consist of neighboring organs, bones, and fat and connective tissue,

the intensities of which should not be considered samples from a single source. While using a single exterior has the advantage of a more stable approximation of appearance, this must be balanced against the lack of positional sensitivity that more local regions provide.

In this paper we explore this balance by experimentally comparing the efficacy in segmentation of image models at three scales. The three models use RIQFs to capture intensity distribution variability. As a baseline for comparison, the first model, “global”, uses two regions, the interior and exterior of the organ of interest, as in [9]. The second model, “local-clustered”, attempts to reflect the inhomogeneity in the intensity distribution around the exterior by determining local distribution-consistent region types relative to the object boundary. The region types are formed by clustering on local RIQFs. We then partition the object boundary according to region type and apply PCA on the cluster populations. For the third model, “local-geometric”, we separately train the intensity distribution and its variability at each object-relative local region.

Section 2 reviews the intensity quantile methodology and describes the three image appearance models. Section 3 describes the segmentation framework and presents results on bladders and prostates in CT in the context of adaptive radiotherapy for prostate cancer. In section 4, we conclude and posit future directions.

2 Quantile Functions and Image Match

We begin with a description of RIQFs, followed by details of the three appearance models. For each regional scale we will construct the image match function.

2.1 Regional Intensity Quantile Functions

[9] describes an approach to image match that probabilistically represents the appearance of an object in an image. The basic unit of appearance is the regional intensity quantile function (RIQF), derived from the intensity histogram within an object-relative region, such as the interior near the object boundary. Quantile functions are a useful parameterization of one-dimensional distributions. For example, RIQFs have the advantage that certain common changes in a distribution, such as mean shift and variance scaling, are represented as linear changes in the RIQF feature space. Given the RIQFs from a region corresponding across image cases, PCA can therefore be used to characterize the variability in that region’s intensity distribution.

We briefly review the construction of the RIQF in the context of the distance metric that provides linearity. Let q and r be the continuous, one-dimensional intensity distributions in two regions between which we wish to measure the similarity. The Mallows distance [10] between q and r , with cumulative distribution functions Q and R , respectively, is defined as

$$M_p(q, r) = \left(\int_0^1 |Q^{-1}(t) - R^{-1}(t)|^p dt \right)^{1/p}. \quad (1)$$

An n -dimensional RIQF is then the discretized inverse cumulative distribution function on intensities in a region, i.e., $Q^{-1}(t)$ or $R^{-1}(t)$ in the above equation. Let these discretized quantile functions be denoted \mathbf{q} or \mathbf{r} . Coordinate j of \mathbf{q} or \mathbf{r} stores the average of the $[\frac{j-1}{n}, \frac{j}{n}]$ quantile of the intensity distribution for that region, i.e., $\mathbf{q}_j = \int_{(j-1)/n}^{j/n} Q^{-1}(t)dt$. After discretization, the Mallows distance above corresponds (up to a scale factor) to the L_p vector norm between \mathbf{q} and \mathbf{r} ,

$$M_p(q, r) \approx \left(\frac{1}{n} \sum_{j=1}^n \|\mathbf{q}_j - \mathbf{r}_j\|^p \right)^{(1/p)}. \quad (2)$$

Through quantile functions, regional intensity distributions are understood as points in an n -dimensional Euclidean space in which distance corresponds to the M_2 metric, and mean and variance changes in intensities are linear.

2.2 Global Regions

In the following sections we describe the construction of the RIQF training populations and the image match used in segmentation, for each of the global, local-clustered, and local-geometric appearance models. The image match itself in each case is a sum over regions of log probabilities in the context of Bayesian deformable model segmentation (sec. 3.1). Our data consist of CT images of the male pelvic region and corresponding manual segmentations of bladders and prostates used for training both the shape and appearance models (sec. 3). Each object is trained separately.

With the global regions model we analyze the intensity patterns near the organ boundary, interior and exterior to it. For each training image I_p , we construct the RIQFs $\mathbf{q}^{in,p}$ and $\mathbf{q}^{out,p}$ through sampling relative to the manual segmentation. The contribution of a voxel is Gaussian weighted by its distance to the surface. The σ for this voxel weighting is a parameter of the training: ours is such that the contribution is negligible farther than 1 cm. We then apply PCA separately to the two RIQF sets, $\{\mathbf{q}^{in,p}, \forall p\}$ and $\{\mathbf{q}^{out,p}, \forall p\}$, to obtain Gaussian models of the intensity variation inside and outside the organ. For segmenting a target case then, the image is similarly sampled relative to a prospective model. We treat the two regions as though they were independent, so the match computed is the sum of the log probabilities of the interior and exterior target RIQFs.

2.3 Local-Clustered Regions

The impetus for the local-clustered appearance model [11] is that more local regions will better specify the exterior than the common single homogeneous region approach. The question is what constitutes a region. For our purposes, consider that an organ or other volume whose local intensity distributions are distinguishable from those of neighboring volumes constitutes such a region.




Fig. 2. Sagittal views of a male pelvis in CT, with bladder boundary colored by region type. Left: a 3D view. Middle: the in-plane contour also colored by region type, with the prostate shown for reference. Right: an off-sagittal 3D view of the same bladder, now with the prostate also in 3D and in the foreground.

Examples of such volumes are neighboring organs and fat deposits. These regions are a cause of intensity inhomogeneity in the organ exterior. We avoid modeling the 3D shapes of these regions and for our appearance model consider them only as they affect the local distributions near the boundary of the object of interest. We determine region types, corresponding to these different organs and volumes, through clustering on the RIQFs of smaller regions.

Over all training images I_p , we compute RIQFs $\mathbf{q}^{i,p}$ for many local exterior regions anchored to individual geometrically defined points on the object boundary indexed by i . The interior of the object is treated as one region. As in the global case, the contribution of a voxel to $\mathbf{q}^{i,p}$ is weighted by its distance to the surface but further is equal to zero if its closest point on the surface is not close enough to point i —ensuring locality. Computing these weights involves minimal additional computation through our sampling scheme (see sec. 3.1). Parameters for this model are the density of points on the surface and the “close enough” distance, in addition to the Gaussian distance weight of the global region model.

We then cluster the pooled set of RIQFs for all boundary points and images, $\{\mathbf{q}^{i,p}, \forall i, p\}$, using Fuzzy C-means Clustering [12]. The Euclidean feature space distances inherent in this method hold for RIQFs (see sec. 2.1). We specify the number of clusters. The results are the set of RIQF cluster centers $\{\mu_k\}$

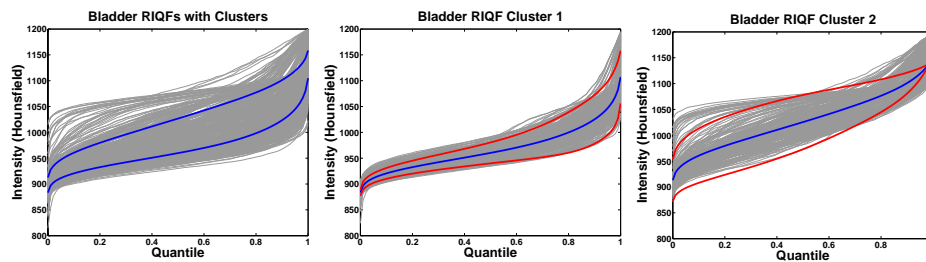


Fig. 3. Clustering on local RIQFs for the bladder exterior. Left shows the pooled data with cluster centers overlaid (two clusters). The middle and right images show the two cluster populations with mean and ± 2 standard deviations overlaid. The reasonable separation into lighter and darker distributions is evidenced in Fig. 2.

and cluster membership scores $\{u_k^{i,p}, \forall i, p\}$ over all positions and images that minimize the sum of score-weighted distance in the M_2 metric used:

$$\{\mu_k\} = \min_{\{\tilde{\mu}_1, \dots, \tilde{\mu}_K \in \mathbb{R}^n\}} \sum_{\forall i, p} \sum_{k=1}^K u_k^{i,p} * \|\mathbf{q}^{i,p} - \tilde{\mu}_k\|^2, \quad (3)$$

with $u_k^{i,p} \in [0, 1]$ and $\sum_k u_k^{i,p} = 1$ for given i and p .

The last two steps for training the local-clustered model are to partition the object boundary according to representative region type and to separately characterize the variability in the cluster populations. To partition the object model boundary, we use the explicit correspondence across training cases required of our shape model (see sec. 3.1). Each point i is assigned the region type/cluster center k that maximizes $\sum_p u_k^{i,p}$, the sum of cluster membership scores for that point over all images. Finally, we perform PCA on the individual cluster populations, resulting in Gaussian models of the RIQF variability per region type. Each point i is then assigned the PCA model of its cluster.

During target segmentation, the image is sampled relative to a prospective model to populate the set of RIQFs $\{\mathbf{q}^i, \forall i\}$. For each point i , we compute the probability of \mathbf{q}^i with respect to the principal modes and variances of that point’s region type. The image match is then the sum of the log probabilities over all points, weighted by the points’ comparative importance (related to how many voxels contributed to each local RIQF).

Confirming evidence for this approach is found in the spatial distribution of region-type on the boundary that we observe in training. When we look for two clusters on the bladder data (Fig. 3), the resultant boundary partition corresponds to lighter and darker local distribution areas (Fig. 2). This is anatomically justified because the bladder is surrounded mostly by lower intensity bowel and fat, with much brighter tissue from the pubic bone area and prostate inferior to it. There is similar evidence for the prostate, which has brighter tissue exterior to it in the pubic bones areas and bladder, with darker tissue elsewhere.

2.4 Local-Geometric Regions

A problem with the global model is that through considering only a single exterior region it sacrifices all positional sensitivity. The appearance model is not specific enough, in that there could be many large-scale regions with the same aggregate tissue mixture. The local-clustered model attempts to capture the exterior inhomogeneity and replaces a single Gaussian model of intensity with a number of local region-type Gaussian models, thus leading to a better specified appearance. However, when we analyze the RIQF data we see that these region-types still over-generalize the local intensity distribution at many particular points on the object boundary. Figure 4 shows a point’s regional data relative to the cluster populations. Neither cluster population is representative of that point’s variability; yet the RIQFs themselves are tightly grouped (at least in that projection). So we must use an even more specific model.

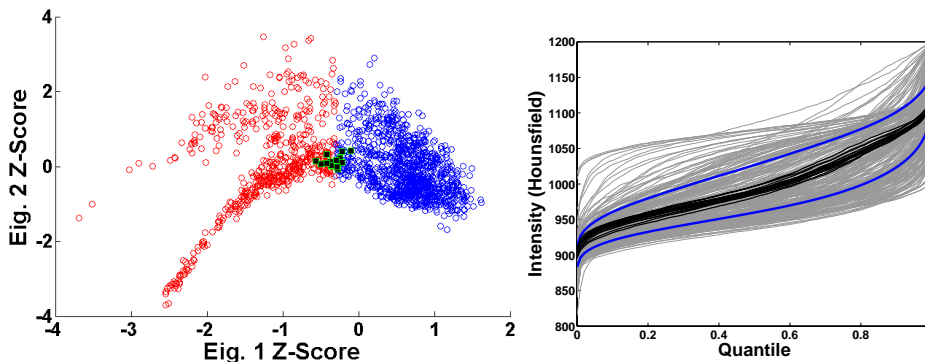


Fig. 4. Plots showing the over-generalization of the local-clustered scheme, for bladder data. Left shows the data projected onto the first two eigen directions of the pooled PCA space, colored according to cluster type (two clusters). Right shows the pooled curve plots, with cluster centers. The highlighted data are the local RIQFs for a particular geometrically corresponding point across training. Note that either cluster population poorly represents that point’s variability.

The local-geometric appearance model addresses the weaknesses of both previous models. In training, the local RIQFs are sampled as in the local-clustered model. However in this scheme, we analyze the variability at each point on its own, applying PCA to the set $\{\mathbf{q}^{i,p}, \forall p\}$, for each point i , thereby constructing a more specific Gaussian model for each local region. At target time then, we treat the local regions as though they were independent. We compute a probability for each point i with respect to the principal modes and variances of that point’s PCA model. The image match is then the weighted sum of log probabilities, as in the previous models.

3 Experimental Results

We compare the efficacy of the three appearance models in the context of deformable model segmentation of bladders and prostates. These organs present a very challenging segmentation problem due to the lack of contrast between the bladder and prostate and the large variability of the bladder across days (Fig. 1). We run the same experiment three times, with the only difference being the image match model. Our data consists of five patient image sets, each of approximately 16 daily CT scans of the male pelvic area taken during radiotherapy courses. The images have an in-plane resolution of 512×512 with voxel dimensions of $0.98 \text{ mm} \times 0.98 \text{ mm}$ and an inter-slice distance of 3 mm. We also are provided expert manual segmentations of the bladder and prostate in every image. We consider the patients separately, segmenting the images from one patient in a leave-one-day-out study, where training is based on all days for the patient except the target day. In section 3.1 we discuss our segmentation framework. In section 3.2 we present segmentation results.

3.1 The Segmentation Framework

In this section we discuss our shape model, its use in segmentation, and the particular parameters of our experiment. We use discrete medial representations, or “m-rep”, models to describe the shape of prostates and bladders [13]. An m-rep is a discretely sampled grid of medial atoms, where each atom consists of a hub and two equal-length spokes. The boundary of the object model passes orthogonally through the spoke ends. Properties, such as spoke length and orientation, are interpolated between atoms in the grid. The model defines a coordinate system which provides an explicit correspondence between deformations of the same m-rep model (e.g., those m-reps making up a training set) and the 3D volume in the object boundary region. This leads to a constructive approach to sampling an image relative to an object, stepping along profiles normal to the surface that are provided by the m-rep.

To extract m-reps from images, we perform Bayesian deformable model segmentation, with a semi-automatic initialization [13]. To start, a mean bladder or prostate model is positioned in a target image using a similarity transform computed from two prostate landmarks. After initialization, we optimize the posterior $p(m|I)$ of the geometric parameters given the image data. This is equivalent to optimizing the sum of the log prior $p(m)$ and the log likelihood $p(I|m)$, which measure geometric typicality and image match, respectively. Geometric typicality and the initial mean models are based on the statistics of m-rep deformation over the training set [14].

There are several parameters that specify the appearance models we use. The choice of region depth has already been discussed. For the clustered and geometric local models, we use a boundary point density that places 306 points on the bladder model surface and 290 on the prostate at fixed object relative coordinates. We set the number of clusters to be two for both object exteriors (see sec. 2.3). Our experiments show that while small changes in these parameters do affect specific results, they do not change the overall conclusions.

3.2 Segmentation Results

We consider the relative segmentation accuracy of the three appearance models by comparing automatically generated results against the expert manual segmentations. As our measures, we use average surface distance (ASD) and volume overlap given by the Dice Similarity Coefficient (DSC) [15], which is intersection over average. We will describe the results of the global model, then the local-clustered model relative to the first, and finally the local-geometric appearance model relative to the first two. Tables 1 and 2 show bladder and prostate results for each patient and appearance model, while Fig. 5 contains trend graphs over the pooled patient data.

The global interior/exterior appearance model results in a median volume overlap of 91.2% for bladders and 92.1% for prostates, with an overlap greater than 90% in about 50 of 80 total target bladders and 60 of 80 prostates. In terms of average surface distance, global regions results in a median ASD of 1.40mm for

bladders and 1.03mm for prostates with 50 of 80 bladders and 65 of 80 prostates having ASD less than 1.5mm. These results are good in the context of the male pelvis in CT, exceeding the agreement we observe between experts.

Segmentations using the local-clustered regions appearance scheme improve upon the global results in 57.5% (46 of 80) of bladders and 53.8% (43 of 80) of prostates over all patients. Considered separately (as they are trained and segmented), this appearance model improves bladder and prostate segmentations in a majority of three of the five patient image sets. In the successful patients, bladders are improved in 68.1% (30 of 44) of images while prostates are improved in 67.3% (33 of 49). These results are encouraging considering that even in the patient sets that are not improved in a majority of the images, the results are not significantly worse (see Tables 1 and 2).

The local-geometric model notably improves results over the first two methods, as evidenced in Fig. 5. In every cumulative measure, this method provides better fits overall. As well, the tables show that in terms of both volume overlap and average surface distance measures, local-geometric is the best method in 4 of 5 patient prostate sets and all 5 patient bladder sets. The improvements are more pronounced in the bladder because there is more room for improvement. The prostate’s mostly rigid day-to-day change is well captured by the initial transform in the segmentation algorithm (sec. 3.1).

4 Conclusions

In this paper, we compared the efficacy of three statistical appearance models based on RIQFs, where the difference in the models is in regional scale. The first model, the common global interior/exterior regions approach, provides good results and is computationally inexpensive. However, it is not a realistic model, in that it assumes the exterior of an organ has no distinguishing features which can

Patient	GlobDSC	ClustDSC	LocalDSC	GlobASD	ClustASD	LocalASD
1	91.0%	92.0	93.1	1.43mm	1.50	1.16
2	93.5	93.6	94.0	1.23	1.15	1.09
3	90.9	91.3	92.8	1.58	1.48	1.21
4	93.7	93.9	95.1	1.16	1.14	0.92
5	89.7	89.9	90.9	2.13	1.98	1.80

Table 1. Bladder median overlap and surface distance per patient. See sec. 3.2 for meaning of the abbreviations.

Patient	GlobDSC	ClustDSC	LocalDSC	GlobASD	ClustASD	LocalASD
1	90.2%	91.8	90.7	0.98mm	0.82	0.93
2	92.0	92.3	94.2	1.34	1.26	0.97
3	92.3	92.0	93.0	0.95	0.94	0.83
4	93.9	94.2	94.2	0.97	0.93	0.90
5	91.3	90.0	91.9	1.59	1.78	1.44

Table 2. Prostate median overlap and surface distance per patient.

be useful to segmentation. The second model, local-clustered, gives a boundary partition consistent with anatomy and leads to somewhat improved results over the global method. However, this model imposes an inaccurate appearance on many particular local regions. The final appearance model, which analyzes the RIQFs per local region, provides the most specific local appearance at every point and results in the best segmentations overall.

Our analysis asserts several directions for continuing research. One area of future research is to control the kind of distribution variability we are attempting to capture in the RIQF framework. PCA and Euclidean clustering on quantile functions is appropriate only to the extent that the variability in the observed distributions is captured by operations that are linear in the RIQF feature space. Some variability, such as changes in the mixture of multiple sources, is non-linear in this space, so it should be accounted for prior to PCA. At the moment when we train with the local-geometric appearance model, we do not explicitly

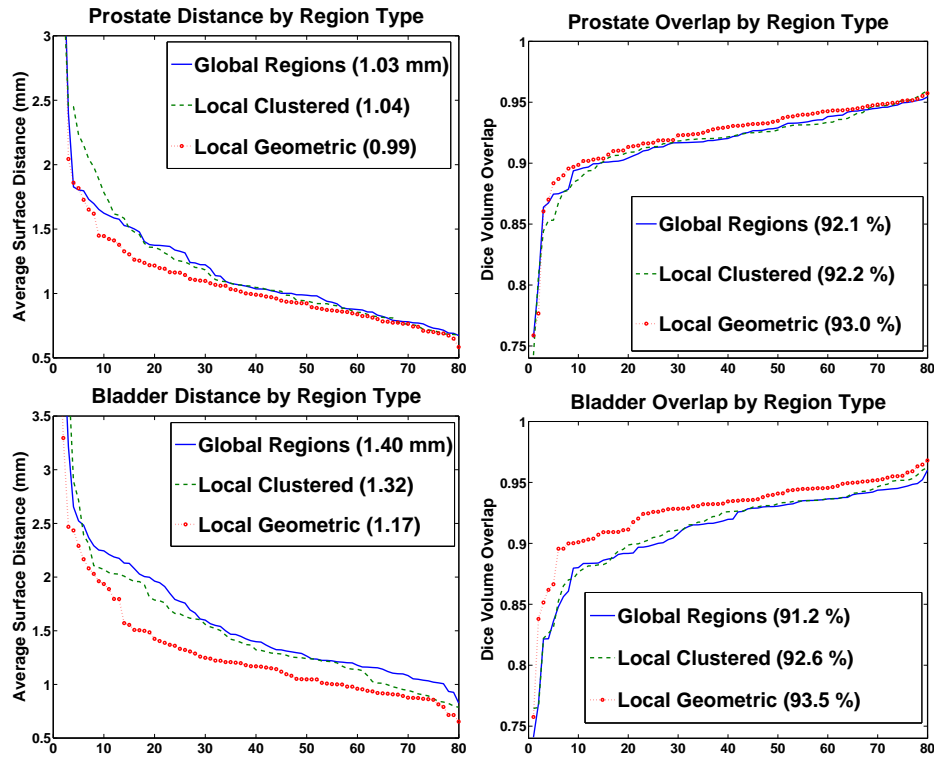


Fig. 5. Comparisons of the three appearance models by organ (prostate and bladder by row), average surface distance (left column) and volume overlap (right column). The comparisons are with respect to expert manual segmentations. The data for each model is sorted independently over all patients to show trends, so the abscissa is image number and does not correspond across models. The numbers in parentheses are median values. While the local-clustered model just outperforms the global, the local-geometric appearance model is clearly superior.

ensure mixture-consistent regions. In the local-clustered model we may get some mixture-consistency in the cluster in that there is distance between observed distributions that differ in mixture.

A key to improving the performance of the local regional methods is correspondence. The geometric correspondence the methods rely on is a baseline for finding what local image intensity distributions to expect, but the *image* correspondence comes from how the neighboring objects are configured with respect to the object of interest. The image correspondence is in general too strongly tied to the geometric correspondence. For example, the region representing the bladder may change position relative to the prostate surface across days, while the geometric correspondence of the prostates will not. The result is a false association of bladder type and non-bladder type distributions based upon geometric correspondence that is a source of confusion in the appearance model. We will address this problem through an extended local-clustered approach, where we model the changes in the region type partition on the object surface.

Other ongoing research involves both the joint modeling of the local RIQFs, which would obviate the assumption of independence in our image match, and a multiscale segmentation scheme that takes advantage of the positional sensitivity the local regions provide. Finally, we look to improve the local-clustered method through better clustering in the space of local RIQFs, such as that proposed by [16]. We would like a clustering that captures the disparate shapes and relative counts of the ideal cluster populations—bone, bowel, fat, etc., rather than only light versus dark RIQFs.

5 Acknowledgements

We are grateful for algorithmic and code contributions from Gregg Tracton and Graham Gash. We thank Joshua Levy for authoring the software pipeline allowing for fast processing of our data set. Finally, we thank the reviewers for their thoughtful comments. This work was funded under NIBIB P01 EB02779.

References

1. Cootes, T.F., Cooper, D., Taylor, C.J., Graham, J.: Active shape models - their training and application. *Computer Vision and Image Understanding* **61**(1) (January 1995) 38–59
2. Stough, J., Pizer, S.M., Chaney, E.L., Rao, M.: Clustering on image boundary regions for deformable model segmentation. In: *IEEE International Symposium on Biomedical Imaging (ISBI)*. (2004)
3. Scott, I.M., Cootes, T.F., Taylor, C.J.: Improving appearance model matching using local image structure. In: *Proc. Int. Conf. on Image Processing in Medical Imaging*. Volume 2732 of *Lecture Notes in Computer Science.*, Springer (2003) 258–269
4. Zhu, S.C., Yuille, A.L.: Region competition: Unifying snakes, region growing, and bayes/MDL for multiband image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **18**(9) (1996) 884–900

5. Tsai, A., Anthony Yezzi, J., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W.E., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging* **22**(2) (2003) 137–154
6. Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Transactions on Image Processing* **10**(2) (2001) 266–277
7. Freedman, D., Radke, R.J., Zhang, T., Jeong, Y., Lovelock, D.M., Chen, G.T.Y.: Model-based segmentation of medical imagery by matching distributions. *IEEE Transactions on Medical Imaging* **24**(3) (March 2005) 281–292
8. Rubner, Y., Puzicha, J., Tomasi, C., Buhmann, J.M.: Empirical evaluation of dissimilarity measures for color and texture. *Computer Vision and Image Understanding* **84** (2001) 25–43
9. Broadhurst, R.E., Stough, J., Pizer, S.M., Chaney, E.L.: A statistical appearance model based on intensity quantile histograms. In: *IEEE International Symposium on Biomedical Imaging (ISBI)*. (2006)
10. Levina, E.: *Statistical Issues in Texture Analysis*. PhD thesis, University of California at Berkeley (2002)
11. Stough, J.V., Broadhurst, R.E., Pizer, S.M., Chaney, E.L.: Clustering on local appearance for deformable model segmentation. In: *IEEE International Symposium on Biomedical Imaging (ISBI)*. (2007)
12. Bezdec, J.C.: *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York (1981)
13. Pizer, S.M., Fletcher, T., Fridman, Y., Fritsch, D.S., Gash, A.G., Glotzer, J.M., Joshi, S., Thall, A., Tracton, G., Yushkevich, P., Chaney, E.L.: Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision - Special UNC-MIDAG issue* **55**(2) (November-December 2003) 85–106
14. Fletcher, P.T., Lu, C., Pizer, S.M., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging* **23**(8) (2004) 995–1005
15. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26** (1945) 297–302
16. Gath, I., Geva, A.: Unsupervised optimal fuzzy clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(7) (1989) 773–780